Journal of Symbolic Computation 47 (2012) 131–152

Contents lists available at SciVerse ScienceDirect

Journal of Symbolic Computation

journal homepage: www.elsevier.com/locate/jsc

Complete subdivision algorithms, II: Isotopic meshing of singular algebraic curves

Michael Burr [a,1], Sung Woo Choi [b], Ben Galehouse [a,2], Chee K. Yap [a,c]

[a] Courant Institute, NYU, 251 Mercer Street, New York, NY 10012, United States

[b] Department of Mathematics, Duksung Women's University, Seoul 132-714, Republic of Korea

[c] Korea Institute of Advanced Study, 85 Hoegiro (Cheongnyangni-dong 207-43), Dongdaemun-gu, Seoul 130-722, Republic of Korea

ARTICLE INFO

Article history:
Received 28 January 2009
Accepted 15 November 2010
Available online 1 September 2011

Keywords:
Meshing
Singularity
Root bound
Evaluation bound
Implicit algebraic curve
Complete numerical algorithm
Subdivision algorithm

ABSTRACT

Given a real valued function $f(X, Y)$, a box region $B_0 \subseteq \mathbb{R}^2$ and $\varepsilon > 0$, we want to compute an ε-isotopic polygonal approximation to the restriction of the curve $S = f^{-1}(0) = \{p \in \mathbb{R}^2 : f(p) = 0\}$ to B_0. We focus on subdivision algorithms because of their adaptive complexity and ease of implementation. Plantinga & Vegter gave a numerical subdivision algorithm that is exact when the curve S is bounded and non-singular. They used a computational model that relied only on function evaluation and interval arithmetic. We generalize their algorithm to any bounded (but possibly non-simply connected) region that does not contain singularities of S. With this generalization as a subroutine, we provide a method to detect isolated algebraic singularities and their branching degree. This appears to be the first complete *purely numerical* method to compute isotopic approximations of algebraic curves with isolated singularities.

1. Introduction

Given $\varepsilon > 0$, a box region $B_0 \subseteq \mathbb{R}^2$ and a real valued function $f : \mathbb{R}^2 \to \mathbb{R}$, we want to compute a polygonal approximation P to the restriction of the implicit curve $S = f^{-1}(0)$ to B_0 (where $f^{-1}(0) =$

E-mail addresses: burr@cims.nyu.edu, mburr1@fordham.edu (M. Burr), swchoi@duksung.ac.kr (S.W. Choi), bgalehou@mpi-inf.mpg.de (B. Galehouse), yap@cs.nyu.edu (C.K. Yap).

[1] Present address: Mathematics Department, Fordham University, Bronx, NY 10458, United States. Tel.: +1 718 817 3298; fax: +1 718 817 3236.

[2] Present address: Max-Planck-Institut für Informtik Department 1: Algorithms and Complexity, D-66123 Saarbrücken, Germany.

0747-7171/$ – see front matter
doi:10.1016/j.jsc.2011.08.021

$\{p \in \mathbb{R}^2 : f(p) = 0\}$). The approximation P must be (1) "topologically correct" and (2) "ε-close" to $S \cap B_0$. We use the standard interpretation of requirement (2), that $d(P, S \cap B_0) \leq \varepsilon$ where $d(\cdot, \cdot)$ is the Hausdorff distance between compact sets. In recent years, it has become accepted (Boissonnat et al., 2006) to interpret requirement (1) to mean that P is isotopic to $S \cap B_0$, which we denote by $P \approx S \cap B_0$. This means that we not only require that P and $S \cap B_0$ are homeomorphic, but also require that they are embedded in $\mathbb{R}^2$ "in the same way". This means that the two embeddings can be continuously deformed to each other, e.g., if $S \cap B_0$ consists of two disjoint ovals, these can be embedded in $\mathbb{R}^2$ as two ovals exterior to each other, or as two nested ovals. Isotopy, but not homeomorphism, requires P to respect this distinction. There is a stronger notion of isotopy called **ambient isotopy** (see the definition in Section 4). We use this stronger notion in this paper (but, for simplicity, we still say "isotopy"). See Boissonnat et al. (2006, p. 183) for a discussion of the connections between ambient and plain isotopy. In this paper, we focus mainly on topological correctness since achieving ε-closeness is not an issue for our particular subdivision approach (cf. Boissonnat et al. (2006, pp. 213–4)). This amounts to setting $\varepsilon = \infty$.

We may call the preceding problem the **2-D implicit meshing problem**. The term "meshing" comes from the corresponding problem in 3-D: given $\varepsilon > 0$ and an implicit surface $S : f(X, Y, Z) = 0$, we want to construct a triangular mesh M such that $d(M, S) \leq \varepsilon$ and $M \approx S$. It is interesting (see Burr et al. (in preparation)) to identify the 1-D meshing with the well-known problem of real root isolation and refinement for a real function $f(X)$.

The **algebraic approach** and the **numerical approach** constitute two extremes of a spectrum among the approaches to most computational problems on curves and surfaces. Algebraic methods can clearly solve most problems in this area, e.g., by an application of the general theory of cylindrical algebraic decomposition (CAD) (Basu et al., 2003). Purely algebraic methods, however, are generally not considered practical, even in the plane (e.g., Hong (1996) and Seidel and Wolpert (2005)), but efficient solutions have been achieved for special cases such as intersecting quadrics in 3-D (Schoemer and Wolpert, 2006). At the other end of the spectrum, the numerical approaches emphasize approximation and iteration. An important class of such algorithms is the class of **subdivision algorithms** which can be viewed as a generalization of binary search. Such algorithms are practical in two senses: they are easy to implement and their complexity is more adaptive with respect to the input instance (Yap, 2006). Another key feature of subdivision algorithms is that they are "localized", meaning that we can restrict our computation to some region of interest.

Besides the algebraic and numerical approaches, there is another approach that might be called the **geometric approach** in which we postulate an abstract computational model with certain (geometric) primitives (e.g., shoot a ray or decide if a point is in a cell). When implementing these geometric algorithms, one must still choose an algebraic or numerical implementation of these primitives. Implementations can also use a hybrid of algebraic and numerical techniques.

Unfortunately, numerical methods seldom have global correctness guarantees. The most famous example is the Marching Cube algorithm (Lorensen and Cline, 1987). Many authors have tried to improve the correctness of subdivision algorithms (e.g., Stander and Hart (1997)). So far, such efforts have succeeded under one of the following situations:

- (A0) Requiring niceness assumptions such as being non-singular or Morse.
- (A1) Invoking algebraic techniques such as resultant computations or manipulations of algebraic numbers.

It is clear that (A0) should be avoided. Generally, we call a method "complete" if the method is correct without any (A0) type restrictions. But many incomplete algorithms (e.g., Marching cube) are quite useful in practice. We want to avoid (A1) conditions because algebraic manipulations are harder to implement and such techniques are relatively expensive and non-adaptive (Yap, 2006). The complete removal of (A0) type restrictions is the major open problem faced by purely numerical approaches to meshing. Thus, Boissonnat et al. (2006, p. 187) state that "*meshing in the vicinity of singularities is a difficult open problem and an active area of research*". Most of the techniques described in their survey are unable to handle singularities. It should be evident that this open problem has an implicit requirement to avoid the use of (A1) techniques.